arXiv:0806.3015v1 [math.OC] 18 Jun 2008

Randomized Methods for Linear Constraints: Convergence Rates and Conditioning

D. Leventhal* A.S. Lewis†

June 18, 2008

Key words: coordinate descent, linear constraint, condition number, randomization, error bound, iterated projections, averaged projections, distance to ill-posedness, metric regularity
AMS 2000 Subject Classification: 15A12, 15A39, 65F10, 90C25

Abstract

We study randomized variants of two classical algorithms: coordinate descent for systems of linear equations and iterated projections for systems of linear inequalities. Expanding on a recent randomized iterated projection algorithm of Strohmer and Vershynin for systems of linear equations, we show that, under appropriate probability distributions, the linear rates of convergence (in expectation) can be bounded in terms of natural linear-algebraic condition numbers for the problems. We relate these condition measures to distances to ill-posedness, and discuss generalizations to convex systems under metric regularity assumptions.

1 Introduction

The condition number of a problem measures the sensitivity of a solution to small perturbations in its input data. For many problems that arise in

*School of Operations Research and Information Engineering, Cornell University, Ithaca, NY 14853, U.S.A. leventhal@orie.cornell.edu

†School of Operations Research and Information Engineering, Cornell University, Ithaca, NY 14853, U.S.A. aslewis@orie.cornell.edu

numerical analysis, there is often a simple relationship between the condition number of a problem instance and the distance to the set of ill-posed problems—those problem instances whose condition numbers are infinite [3]. For example, with respect to the problem of inverting a matrix A, it is known (see [11], for example) that if A is perturbed to $A + E$ for sufficiently small E, then

$$\frac{\|(A+E)^{-1} - A^{-1}\|}{\|A^{-1}\|} \leq \|A^{-1}\|\|E\| + O(\|E\|^2).$$

Thus, a condition measure for this problem may be taken as $\|A^{-1}\|$. Associated with this is the classical Eckart-Young theorem found in [7], relating the above condition measure to the distance to ill-posedness.

Theorem 1.1 (Eckart-Young) *For any non-singular matrix, A,*

$$\min_E\{\|E\| : A + E \text{ is singular}\} \;=\; \frac{1}{\|A^{-1}\|}.$$

We are typically concerned with relative condition numbers as introduced by Demmel in [3]. For example, with respect to the problem of matrix inversion, the relative condition number is $k(A) := \|A\|\|A^{-1}\|$, the commonly used condition measure.

Condition numbers are also important from an algorithmic perspective. In the above example of matrix inversion, for instance, the sensitivity of a problem under perturbations could come into prominence regarding errors in either the initial problem data or accumulated computational error due to rounding. Hence, it seems natural that condition numbers should affect algorithm speed. For example, in the context of linear programming, Renegar defined a condition measure based on the distance to ill-posedness in [19]– in a similar sense as the Eckart-Young result–and showed its effect on the convergence rate of interior point methods in [20].

For another example, consider the problem of finding a solution to the system $Ax = b$ where A is a positive-definite matrix. It was shown in [1] that the steepest descent method is linearly convergent with rate $(\frac{k(A)-1}{k(A)+1})^2$ and that this bound is asymptotically tight for almost all choices of initial iterates. Similarly, it is well known (see [8]) that the conjugate gradient method applied to the same problem is also linearly convergent with rate $\frac{\sqrt{k(A)}-1}{\sqrt{k(A)}+1}$.

From a computational perspective, a related and important area of study is that of error bounds. Given a subset of a Euclidean space, an error bound is an inequality that bounds the distance from a test vector to the specified subset in terms of some residual function that is typically easy to compute. In that sense, an error bound can be used both as part of a stopping rule during implementation of an algorithm as well as an aide in proving algorithmic convergence. A comprehensive survey of error bounds for a variety of problems arising in optimization can be found in [15].

With regards to the problem of solving a nonsingular linear system $Ax = b$, one connection between condition measures and error bounds is immediate. Let x^* be a solution to the system and x be any other vector. Then

$$\|x - x^*\| = \|A^{-1}A(x - x^*)\| = \|A^{-1}(Ax - b)\| \leq \|A^{-1}\|\|Ax - b\|,$$

so the distance to the solution set is bounded by a constant multiple of the residual vector, $\|Ax - b\|$, and this constant is the same one that appears in the context of conditioning and distance to infeasibility. As we discuss later, this result is not confined to systems of linear equations.

As a result, error bounds and the related condition numbers often make a prominent appearance in convergence proofs for a variety of algorithms. In this paper, motivated by a recent randomized iterated projection scheme for systems of linear equations due to Strohmer and Vershynin in [22], we revisit some classical algorithms and show that, with an appropriate randomization scheme, we can demonstrate convergence rates directly in terms of these natural condition measures. The rest of the paper is organized as follows. In the remainder of this section, we define some notation used throughout the rest of this paper. In Section 3, we consider the problem of solving a linear system $Ax = b$ and show that a randomized coordinate descent scheme, implemented according to a specific probability distribution, is linearly convergent with a rate expressible in terms of traditional conditioning measures. In Section 4, we build upon the work of Strohmer and Vershynin in [22] by considering randomized iterated projection algorithms for linear inequality systems. In particular, we show how randomization can provide convergence rates in terms of the traditional Hoffman error bound in [10] as well as in terms of Renegar's distance to infeasibility from [18]. In Section 5, we consider randomized iterated projection algorithms for general convex sets and, under appropriate metric regularity assumptions, obtain *local* convergence rates in terms of the modulus of regularity.

The classical, deterministic versions of the simple algorithms we consider here have been widely studied, in part due to the extreme simplicity of each iteration: their linear convergence is well-known. However, as remarked for linear systems of equations in [22], randomized versions are interesting for several reasons. The randomized iterated projection method for linear equations from which this work originated may have some practical promise, even compared with conjugate gradients, for example [22]. Our emphasis here, however, is theoretical: randomization here provides a framework for simplifying the analysis of algorithms, allowing easy bounds on the rates of linear convergence in terms of natural linear-algebraic condition measures, such as relative condition numbers, Hoffman constants, and the modulus of metric regularity.

2 Notation

On the Euclidean space $\mathbf{R}^n$, we denote the Euclidean norm by $\|\cdot\|$. Let e_i denote the column vector with a 1 in the i^{th} position and zeros elsewhere.

We consider m-by-n real matrices A. We denote the set of rows of A by $\{a_1^T, \ldots, a_m^T\}$ and the set of columns is denoted $\{A_1, \ldots, A_n\}$. The *spectral norm* of A is the quantity $\|A\|_2 := \max_{\|x\|=1} \|Ax\|$ and the *Frobenius norm* is $\|A\|_F := \sum_{i,j} a_{ij}^2$. These norms satisfy the following inequality [11]:

$$\|A\|_F \leq \sqrt{n}\|A\|_2. \tag{2.1}$$

For an arbitrary matrix, A, let $\|A^{-1}\|_2$ be the smallest constant M such that $\|Ax\|_2 \geq \frac{1}{M}\|x\|_2$ for all vectors x. In the case $m \geq n$, if A has singular values $\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_n$, then M can also be expressed as the reciprocal of the minimum singular value σ_n, and, if A is invertible, this quantity equals the spectral norm of A^{-1}.

The *relative condition number* of A is the quantity $k(A) := \|A\|_2\|A^{-1}\|_2$; related to this is the *scaled condition number* introduced by Demmel in [4], defined by $\kappa(A) := \|A\|_F\|A^{-1}\|_2$. From this, it is easy to verify (using the singular value decomposition, for example) the following relationship between condition numbers:

$$1 \leq \frac{\kappa(A)}{\sqrt{n}} \leq k(A).$$

Now suppose the matrix A is n-by-n symmetric and positive definite. The *energy norm* (or A-norm), denoted $\|\cdot\|_A$, is defined by $\|x\|_A := \sqrt{x^T Ax}$. The

inequality

$$\|x\|_A^2 \leq \|A^{-1}\|_2 \cdot \|Ax\|^2 \quad \text{for all } x \in \mathbf{R}^n \tag{2.2}$$

is useful later. Further, if A is simply positive semi-definite, we can generalize Inequality 2.2:

$$x^T Ax \leq \frac{1}{\underline{\lambda}(A)} \|Ax\|^2 \tag{2.3}$$

where $\underline{\lambda}(A)$ is the smallest non-zero eigenvalue of A. We denote the trace of A by $\operatorname{tr} A$: it satisfies the inequality

$$\|A\|_F \geq \frac{\operatorname{tr} A}{\sqrt{n}}. \tag{2.4}$$

Given a nonempty closed convex set S, let $P_S(x)$ be the projection of x onto S: that is, $P_S(x)$ is the vector y that is the optimal solution to $\min_{z \in S} \|x - z\|_2$. Additionally, define the distance from x to a set S by

$$d(x, S) = \min_{z \in S} \|x - z\|_2 = \|x - P_S(x)\|.$$

The following useful inequality is standard:

$$\|y - x\|^2 - \|P_S(y) - x\|^2 \geq \|y - P_S(y)\|^2 \quad \text{for all } x \in S,\ y \in \mathbf{R}^n. \tag{2.5}$$

3 Randomized Coordinate Descent

Let A be an n-by-n positive-definite matrix. We consider a linear system of the form $Ax = b$, with solution $x^* = A^{-1}b$. We consider the equivalent problem of minimizing the strictly convex quadratic function

$$f(x) = \frac{1}{2} x^T Ax - b^T x,$$

and note the standard relationship

$$f(x) - f(x^*) = \frac{1}{2} \|x - x^*\|_A^2. \tag{3.1}$$

Suppose our current iterate is x and we obtain a new iterate x_+ by performing an exact line search in the nonzero direction d: that is, x_+ is the solution to $\min_{t \in \mathbf{R}} f(x + td)$. This gives us

$$x_+ = x + \frac{(b - Ax)^T d}{d^T Ad} d$$

and

$$(3.2)\quad f(x_+) - f(x^*) = \frac{1}{2}\|x_+ - x^*\|_A^2 = \frac{1}{2}\|x - x^*\|_A^2 - \frac{((Ax-b)^T d)^2}{2d^T Ad}.$$

One natural choice of a set of easily-computable search directions is to choose d from the set of coordinate directions, $\{e_1, \ldots, e_n\}$. Note that, when using search direction e_i, we can compute the new point

$$x_+ = x + \frac{b_i - a_i^T x}{a_{ii}} e_i$$

using only $2n + 2$ arithmetic operations. If the search direction is chosen at each iteration by successively cycling through the set of coordinate directions, then the algorithm is known to be linearly convergent but with a rate not easily expressible in terms of typical matrix quantities (see [8] or [17]) . However, by choosing a coordinate direction as a search direction randomly according to an appropriate probability distribution, we can obtain a convergence rate in terms of the relative condition number. This is expressed in the following result.

Algorithm 3.3 *Consider an n-by-n positive semidefinite system $Ax = b$ and let $x_0 \in \mathbf{R}^n$ be an arbitrary starting point. For $j = 0, 1, \ldots,$ compute*

$$x_{j+1} = x_j + \frac{b_i - a_i^T x_j}{a_{ii}} e_i$$

where, at each iteration j, the index i is chosen independently at random from the set $\{1, \ldots, n\}$, with distribution

$$P\{i = k\} \;=\; \frac{a_{kk}}{\operatorname{tr} A}.$$

Notice in the algorithm that the matrix A may be singular, but that nonetheless $a_{ii} > 0$ almost surely. If A is merely positive semidefinite, solutions of the system $Ax = b$ coincide with minimizers of the function f, and consistency of the system is equivalent to f being bounded below. We now have the following result.

Theorem 3.4 *Consider a consistent positive-semidefinite system $Ax = b$, and define the corresponding objective and error by*

$$\begin{aligned} f(x) &= \frac{1}{2}x^T Ax - b^T x \\ \delta(x) &= f(x) - \min f. \end{aligned}$$

Then Algorithm 3.3 is linearly convergent in expectation: indeed, for each iteration $j = 0, 1, 2, \ldots$,

$$E[\delta(x_{j+1}) \mid x_j] \;\leq\; \Big(1 - \frac{\underline{\lambda}(A)}{\operatorname{tr} A}\Big)\delta(x_j).$$

In particular, if A is positive-definite and $x^* = A^{-1}b$, we have the equivalent property

$$E[\|x_{j+1} - x^*\|_A^2 \mid x_j] \;\leq\; \Big(1 - \frac{1}{\|A^{-1}\|_2 \operatorname{tr} A}\Big)\|x_j - x^*\|_A^2.$$

Hence the expected reduction in the squared error $\|x_j - x^*\|_A^2$ is at least a factor

$$1 - \frac{1}{\sqrt{n}\kappa(A)} \;\leq\; 1 - \frac{1}{nk(A)}$$

at each iteration.

Proof Note that if coordinate direction e_i is chosen during iteration j, then Equation 3.2 shows

$$f(x_{j+1}) = f(x_j) - \frac{(b_i - a_i^T x_j)^2}{2a_{ii}}.$$

Hence, it follows that

$$E[f(x_{j+1}) \mid x_j] = f(x_j) - \sum_{i=1}^{n} \frac{a_{ii}}{\operatorname{tr}(A)} \frac{(b_i - a_i^T x_j)^2}{2a_{ii}} = f(x_j) - \frac{1}{2\operatorname{tr} A}\|Ax_j - b\|^2.$$

Using Inequality 2.3 and Equation 3.1, we easily verify

$$\frac{1}{2}\|Ax_j - b\|^2 \geq \underline{\lambda}(A)\delta(x_j),$$

and the first result follows. Applying Equation 3.1 provides the second result. The final result comes from applying Inequalities 2.1 and 2.4. □

The simple idea behind the proof of Theorem 3.4 is the main engine driving the remaining results in this paper. Fundamentally, the idea is to choose a probability distribution so that the expected distance to the solution from the new iterate is the distance to the solution from the old iterate minus some multiple of a residual. Then, using some type of error bound to bound the distance to a solution in terms of the residual, we obtain expected linear convergence of the algorithm.

Now let us consider the more general problem of finding a solution to a linear system $Ax = b$ where A is an $m \times n$. More generally, since the system might be inconsistent, we seek a "least squares solution" by minimizing the function $\|Ax - b\|^2$. The minimizers are exactly the solutions of the positive-semidefinite system $A^T Ax = A^T b$, to which we could easily apply the previous algorithm; however, as usual, we wish to avoid computing the new matrix $A^T A$ explicitly. Instead, we can proceed as follows.

Algorithm 3.5 *Consider a linear system $Ax = b$ for a nonzero m-by-n matrix A. Let $x_0 \in \mathbf{R}^n$ be an arbitrary initial point and let $r_0 = b - Ax_0$ be the initial residual. For each $j = 0, 1, \ldots,$ compute*

$$\begin{aligned} \alpha_j &= \frac{A_i^T r_j}{\|A_i\|^2} \\ x_{j+1} &= x_j + \alpha_j e_i \\ r_{j+1} &= r_j - \alpha_j A_i, \end{aligned}$$

where, at each iteration j, the index i is chosen independently at random from the set $\{1, \ldots, n\}$, with distribution

$$P\{i = k\} = \frac{\|A_k\|^2}{\|A\|_F^2} \quad (k = 1, 2, \ldots, n).$$

(In the formula for α_j, notice by assumption that $A_i \neq 0$ almost surely.)

Note that the step size at each iteration can be obtained by directly minimizing the residual in the respective coordinate direction. However, the algorithm can also be viewed as the application of the algorithm for positive definite systems on the system of normal equations, $A^T Ax = A^T b$, without actually having to compute the matrix $A^T A$. Given the motivation of directly minimizing the residual, we would expect that Algorithm 3.5 would converge to a least squares solution, even in the case where the underlying system is inconsistent. The next result shows that this is, in fact, the case.

Theorem 3.6 *Consider any linear system $Ax = b$, where the matrix A is nonzero. Define the least-squares residual and the error by*

$$\begin{aligned} f(x) &= \frac{1}{2}\|Ax - b\|^2 \\ \delta(x) &= f(x) - \min f. \end{aligned}$$

Then Algorithm 3.5 is linearly convergent in expectation to a least squares solution for the system: for each iteration $j = 0, 1, 2 \ldots$,

$$E[\delta(x_{j+1}) \mid x_j] \;\leq\; \Big(1 - \frac{\underline{\lambda}(A^T A)}{\|A\|_F^2}\Big)\delta(x_j).$$

In particular, if A has full column rank, we have the equivalent property

$$E[\|x_{j+1} - \hat{x}\|_{A^T A}^2 \mid x_j] \;\leq\; \Big(1 - \frac{1}{\kappa(A)^2}\Big)\|x_j - \hat{x}\|_{A^T A}^2$$

where $\hat{x} = (A^T A)^{-1} A^T b$ is the unique least-squares solution.

Proof It is easy to verify, by induction on j, that the iterates x_j are exactly the same as the iterates generated by Algorithm 3.3, when applied to the positive semi-definite system $A^T Ax = A^T b$, and furthermore that the residuals satisfy $r_j = b - Ax_j$ for all $j = 0, 1, 2, \ldots$. Hence, the results follow directly by Theorem 3.4. □

By the coordinate descent nature of this algorithm, once we have computed the initial residual r_0 and column norms $\{\|A_i\|^2\}_{i=1}^n$, we can perform each iteration in $O(n)$ time, just as in the positive-definite case. Specifically, this new iteration takes $4n + 1$ arithmetic operations, compared with $2n + 2$ for the positive-definite case.

For a computational example, we apply Algorithm 3.5 to random $500 \times n$ matrices where each element of A and b is an independent Gaussian random variable and we let n take values 50, 100, 150 and 200.

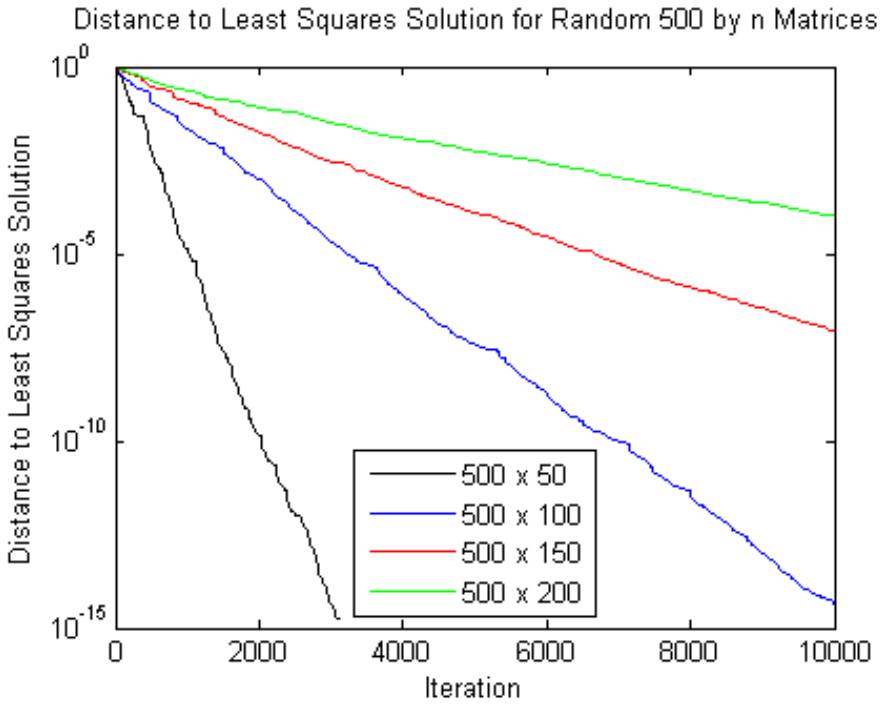

n	Theoretical Rate	Observed Rate
50	0.98995	0.98911
100	0.99681	0.99670
150	0.99855	0.99837
200	0.99928	0.99908

Note that in the above examples, the theoretical bound provided by Theorem 3.6 predicts the actual behavior of the algorithm reasonably well.

4 Randomized Iterated Projections

Iterated projection algorithms share some important characteristics with coordinate descent algorithms. Both are well studied and much convergence theory exists; a comprehensive overview on iterated projections can be found in [5]. However, even for linear systems of equations, standard developments do not provide bounds on convergence rates in terms of usual condition numbers. By contrast, in the recent paper [22], Strohmer and Vershynin obtained such bounds via the following randomized iterated projection algorithm, which also provided the motivation for our work in the previous section.

Algorithm 4.1 *Consider a linear system $Ax = b$ for a nonzero m-by-n matrix A. Let $x_0 \in \mathbf{R}^n$ be an arbitrary initial point. For each $j = 0, 1, \ldots,$ compute*

$$x_{j+1} = x_j - \frac{a_i^T x_j - b_i}{\|a_i\|^2} a_i$$

where, at each iteration j, the index i is chosen independently at random from the set $\{1, \ldots, m\}$, with distribution

$$P\{i = k\} \;=\; \frac{\|a_k\|^2}{\|A\|_F^2} \quad (k = 1, 2, \ldots, m).$$

Notice that the new iterate x_{j+1} is simply the orthogonal projection of the old iterate x_j onto the hyperplane $\{x : a_i^T x = b_i\}$. At first sight, the choice of probability distribution may seem curious, since we could rescale the equations arbitrarily without having any impact on the projection operations. However, following [22], we emphasize that the aim is to understand linear convergence rates in terms of *linear-algebraic* condition measures associated with the original system, rather than in terms of *geometric* notions associated with the hyperplanes. This randomized algorithm has the following behavior.

Theorem 4.2 (Strohmer-Vershynin, [22]) *Given any matrix A with full column rank, suppose the linear system $Ax = b$ has solution x^*. Then Algorithm 4.1 converges linearly in expectation: for each iteration $j = 0, 1, 2, \ldots,$*

$$E[\|x_{j+1} - x^*\|_2^2 \mid x_j] \;\leq\; \Big(1 - \frac{1}{\kappa(A)^2}\Big)\|x_j - x^*\|_2^2$$

We seek a way of generalizing the above algorithm and convergence result to more general systems of linear inequalities, of the form

$$\left\{ \begin{array}{rcll} a_i^T x & \leq & b_i & (i \in I_{\leq}) \\ a_i^T x & = & b_i & (i \in I_{=}), \end{array} \right. \tag{4.3}$$

where the disjoint index sets $I_{\leq}$ and $I_{=}$ partition the set $\{1, 2, \ldots, m\}$. To do so, staying with the techniques of the previous section, we need a corresponding error bound for a system of linear inequalities. First, given a vector $x \in \mathbf{R}^n$, define the vector x^+ by $(x^+)_i = \max\{x_i, 0\}$. Then a starting point for this subject is a result by Hoffman in [10].

Theorem 4.4 (Hoffman) *For any right-hand side vector $b \in \mathbf{R}^m$, let S_b be the set of feasible solutions of the linear system (4.3). Then there exists a constant L, independent of b, with the following property:*

$$x \in \mathbf{R}^n \text{ and } S_b \neq \emptyset \quad \Rightarrow \quad d(x, S_b) \leq L\|e(Ax-b)\|, \tag{4.5}$$

where the function $e\colon \mathbf{R}^m \to \mathbf{R}^m$ is defined by

$$e(y)_i = \begin{cases} y_i^+ & (i \in I_{\leq}) \\ y_i & (i \in I_{=}), \end{cases}$$

In the above result, each component of the vector $e(Ax-b)$ indicates the error in the corresponding inequality or equation. In particular $e(Ax-b) = 0$ if and only if $x \in S_b$. Thus Hoffman's result provides a linear bound for the distance from a trial point x to the feasible region in terms of the size of the "a posteriori error" associated with x.

We call the minimum constant L such that property (4.5) holds the *Hoffman constant* for the system (4.3). Several authors give geometric or algebraic meaning to this constant, or exact expressions for it, including [9], [14], [13]; for a more thorough treatment of the subject, see [15]. In the case of linear equations (that is, $I_{\leq} = \emptyset$), an easy calculation using the singular value decomposition shows that the Hoffman constant is just the reciprocal of the smallest nonzero singular value of the matrix A, and hence equals $\|A^{-1}\|_2$ when A has full column rank.

For the problem of finding a solution to a system of linear inequalities, we consider a randomized algorithm generalizing Algorithm 4.1.

Algorithm 4.6 *Consider the system of inequalities (4.3). Let x_0 be an arbitrary initial point. For each $j = 0, 1, \ldots,$ compute*

$$\begin{aligned} \beta_j &= \begin{cases} (a_i^T x_j - b_i)^+ & (i \in I_{\leq}) \\ a_i^T x_j - b_i & (i \in I_{=}) \end{cases} \\ x_{j+1} &= x_j - \frac{\beta_j}{\|a_i\|^2} a_i \end{aligned}$$

where, at each iteration j, the index i is chosen independently at random from the set $\{1, \ldots, m\}$, with distribution

$$P\{i = k\} = \frac{\|a_k\|^2}{\|A\|_F^2} \quad (k = 1, 2, \ldots, m).$$

In the above algorithm, notice $\beta_j = e(Ax_j - b)_i$. We can now generalize Theorem 4.2 as follows.

Theorem 4.7 *Suppose the system (4.3) has nonempty feasible region S. Then Algorithm 4.6 converges linearly in expectation: for each iteration $j = 0, 1, 2, \ldots$,*

$$E[d(x_{j+1}, S)^2 \mid x_j] \;\leq\; \Big(1 - \frac{1}{L^2 \|A\|_F^2}\Big) d(x_j, S)^2.$$

where L is the Hoffman constant.

Proof Note that if the index i is chosen during iteration j, then it follows that

$$\begin{aligned}
\|x_{j+1} - P_S(x_{j+1})\|_2^2 \;&\leq\; \|x_{j+1} - P_S(x_j)\|_2^2 \\
&=\; \Big\|x_j - \frac{e(Ax_j - b)_i}{\|a_i\|^2} a_i - P_S(x_j)\Big\|_2^2 \\
&=\; \|x_j - P_S(x_j)\|_2^2 + \frac{e(Ax_j - b)_i^2}{\|a_i\|^2} - 2\frac{e(Ax_j - b)_i}{\|a_i\|^2} a_i^T (x_j - P_S(x_j)).
\end{aligned}$$

Note $P_S(x_j) \in S$. Hence if $i \in I_{\leq}$, then $a_i^T P_S(x_j) \leq b_i$, and $e(Ax_j - b)_i \geq 0$, so

$$e(Ax_j - b)_i a_i^T (x_j - P_S(x_j)) \geq e(Ax_j - b)_i (a_i^T x_j - b_i) = e(Ax_j - b)_i^2.$$

On the other hand, if $i \in I_{=}$, then $a_i^T P_S(x_j) = b_i$, so

$$e(Ax_j - b)_i a_i^T (x_j - P_S(x_j)) = e(Ax_j - b)_i (a_i^T x_j - b_i) = e(Ax_j - b)_i^2.$$

Putting these two cases together with the previous inequality shows

$$d(x_{j+1}, S)^2 \;\leq\; d(x_j, S)^2 - \frac{e(Ax_j - b)_i^2}{\|a_i\|^2}.$$

Taking the expectation with respect to the specified probability distribution, it follows that

$$E[d(x_{j+1}, S)^2 \mid x_j] \leq d(x_j, S)^2 - \frac{\|e(Ax_j - b)\|^2}{\|A\|_F^2}$$

and the result now follows by the Hoffman bound. □

Since Hoffman's bound is not independent of the scaling of the matrix A, it is not surprising that a normalizing constant like $\|A\|_F^2$ term appears in the result.

For a computational example, we consider linear inequality systems $Ax \leq b$ where the elements of A are independent standard Gaussian random variables and b is chosen so that the resulting system has a non-empty interior. We consider matrices A which are $500 \times n$, letting n take values 50, 100, 150 and 200. We then apply Algorithm 4.6 to these problems and observe the following computational results.

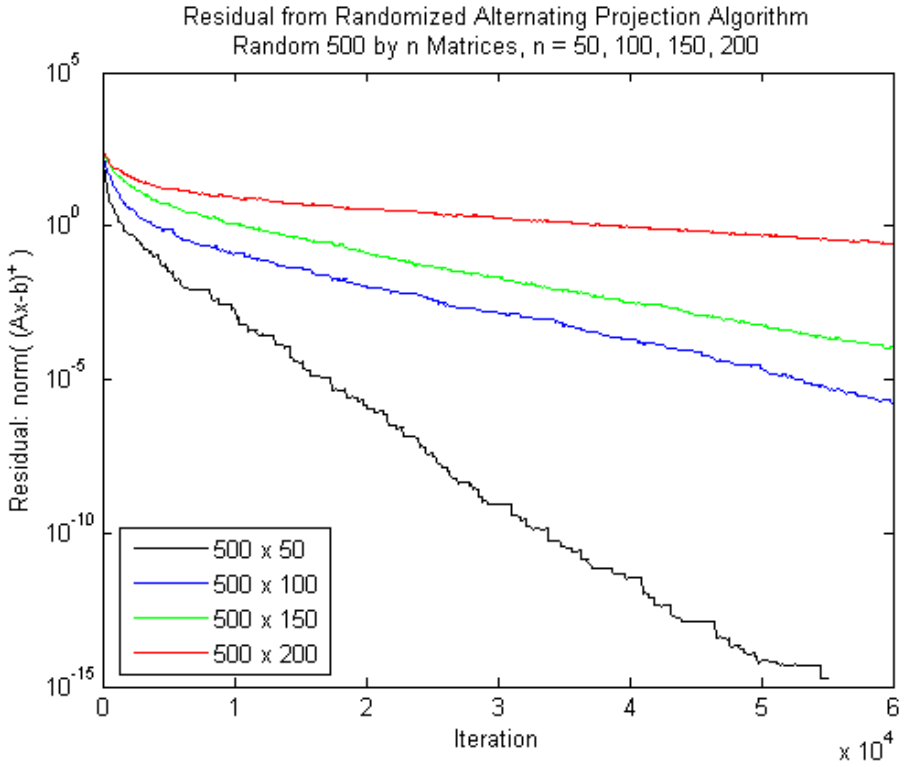

Another natural conditioning measure for linear inequality systems is the distance to infeasibility, defined by Renegar in [18], and shown in [20] to govern the convergence rate of interior point methods for linear programming. It is interesting, therefore, from a theoretical perspective, to obtain a linear convergence rate for iterated projection algorithms in terms of this condition

measure as well. For simplicity, we concentrate on the inequality case, $Ax \leq b$. To begin, let us recall the following results.

The *distance to infeasibility* [18] for the system $Ax \leq b$ is the number

$$\mu \;=\; \inf\Big\{ \max\{\|\Delta A\|, \|\Delta b\|\} : (A + \Delta A)x \leq b + \Delta b \text{ is infeasible}\Big\}.$$

Theorem 4.8 (Renegar, [18], Thm 1.1) *Suppose $\mu > 0$. Then there exists a point $\hat{x} \in S$ satisfying $\|\hat{x}\| \leq \|b\|/\mu$. Furthermore, any point $x \in \mathbf{R}^n$ satisfies the inequality*

$$d(x, S) \;\leq\; \frac{\max\{1, \|x\|\}}{\mu} \|(Ax - b)^+\|.$$

Using this, we can bound the linear convergence rate for the Algorithm 4.6 in terms of the distance to infeasibility, as follows. Notice first that $\|x_j - \hat{x}\|$ is nonincreasing in j, by Inequality 2.5. Suppose we start Algorithm 4.6 at the initial point $x_0 = 0$. Applying Theorem 4.8, we see that for all $j = 1, 2, \ldots,$

$$\|x_j\| \leq \|\hat{x}\| + \|x_j - \hat{x}\| \leq \|\hat{x}\| + \|x_0 - \hat{x}\| \leq \frac{2\|b\|}{\mu},$$

so

$$d(x_j, S) \leq \max\Big\{\frac{1}{\mu}, \frac{2\|b\|}{\mu^2}\Big\} \|(Ax_j - b)^+\|.$$

Using this inequality in place of Hoffman's bound in the proof of Theorem 4.7 gives

$$E[d(x_{j+1}, S)^2 \mid x_j] \leq \left[1 - \frac{1}{\|A\|_F^2 (\max\{\frac{1}{\mu}, \frac{2\|b\|}{\mu^2}\})^2}\right] d(x_j, S)^2.$$

Although this bound may not be the best possible (and, in fact, it may not be as good as the bound provided in Theorem 4.7), this result simply emphasizes a relationship between algorithm speed and conditioning measures that appears naturally in other contexts. In the next section, we proceed with these ideas in a more general framework.

5 Metric Regularity and Local Convergence

The previous section concerned global rates of linear convergence. If instead we are interested in *local* rates, we can re-examine a generalization of our

problem through an alternative perspective of set-valued mappings. Consider a set-valued mapping $\Phi : \mathbf{R}^n \rightrightarrows \mathbf{R}^m$ and the problem of solving the associated constraint system of the form $b \in \Phi(x)$ for the unknown vector x. For example, finding a feasible solution to $Ax \leq b$ is equivalent to finding an x such that

$$b \in Ax + \mathbf{R}^m_+. \tag{5.1}$$

Related to this is the idea of *metric regularity* of set-valued mappings. We say the set-valued mapping Φ is metrically regular at $\bar{x}$ for $\bar{b} \in \Phi(\bar{x})$ if there exists $\gamma > 0$ such that

$$d(x, \Phi^{-1}(b)) \leq \gamma d(b, \Phi(x)) \text{ for all } (x, b) \text{ near } (\bar{x}, \bar{b}), \tag{5.2}$$

where $\Phi^{-1}(b) = \{x : b \in \Phi(x)\}$. Further, the *modulus of regularity* is the infimum of all constants γ such that Equation 5.2 holds. Metric regularity is strongly connected with a variety of ideas from variational analysis: a good background reference is [21].

Metric regularity generalizes the error bounds discussed in previous sections at the expense of only guaranteeing a bound in local terms. For example, if Φ is a single-valued linear map, then the modulus of regularity (at any $\bar{x}$ for any $\bar{b}$) corresponds to the typical conditioning measure $\|\Phi^{-1}\|$ (with $\|\Phi^{-1}\| = \infty$ implying the map is not metrically regular) and if Φ is a smooth single-valued mapping, then the modulus of regularity is the reciprocal of the minimum singular value of the Jacobian, $\nabla\Phi(x)$. From an alternative perspective, metric regularity provides a framework for generalizing the Eckart-Young result on the distance to ill-posedness of linear mappings cited in Theorem 1.1. Specifically, if we define the *radius of metric regularity* at $\bar{x}$ for $\bar{b}$ for a set-valued mapping Φ between finite dimensional spaces by

$$\text{rad}\Phi(\bar{x}|\bar{b}) \;=\; \inf\{\|E\| : \Phi + E \text{ not metrically regular at } \bar{x} \text{ for } \bar{b} + E(\bar{x})\},$$

where the infimum is over all linear functions E, then one obtains the strikingly simple relationship (see [6])

$$\text{modulus of regularity of } \Phi \text{ at } \bar{x} \text{ for } \bar{b} \;=\; \frac{1}{\text{rad}\Phi(\bar{x}|\bar{b})}.$$

We will not be directly using the above result. Here, we simply use the fundamental idea of metric regularity which says that the distance from a point to the solution set, $d(x, \Phi^{-1}(b))$, is locally bounded by some constant

times a "residual". For example, in the case where Φ corresponds to the linear inequality system (5.1), we have that $d(b, \Phi(x)) = \|(Ax - b)^+\|$ implies that the modulus of regularity is in fact a global bound and equals the Hoffman bound. More generally, we wish to emphasize that metric regularity ties together several of the ideas from previous sections at the expense of those results now only holding locally instead of globally.

In what follows, assume all distances are Euclidean distances. We wish to consider how the modulus of regularity of Φ affects the convergence rate of iterated projection algorithms. We remark that linear convergence for iterated projection methods on convex sets has been very widely studied: see [5], for example. Our aim here is to observe, by analogy with previous sections, how randomization makes the linear convergence rate easy to interpret in terms of metric regularity.

Let $S_1, S_2, \ldots, S_m$ be closed convex sets in a Euclidean space $\mathbf{E}$ such that $\cap_i S_i \neq \emptyset$. Then, in a manner similar to [12], we can endow the product space $\mathbf{E}^m$ with the inner product

$$\langle (u_1, u_2, \ldots, u_m), (v_1, v_2, \ldots, v_m) \rangle = \sum_{i=1}^{m} \langle u_i, v_i \rangle$$

and consider the set-valued mapping $\Phi : \mathbf{E} \to \mathbf{E}^m$ given by

$$\Phi(x) = (S_1 - x, S_2 - x, \ldots, S_m - x). \tag{5.3}$$

Then it clearly follows that $\bar{x} \in \cap_i S_i \Leftrightarrow 0 \in \Phi(\bar{x})$. Under appropriate regularity assumptions, we obtain the following local convergence result.

Theorem 5.4 *Suppose the set-valued mapping Φ given by Equation 5.3 is metrically regular at $\bar{x}$ for 0 with regularity modulus γ. Let $\bar{\gamma}$ be any constant strictly larger than γ and let x_0 be any initial point sufficiently close to $\bar{x}$. Further, suppose that $x_{j+1} = P_{S_i}(x_j)$ with probability $\frac{1}{m}$ for $i = 1, \ldots, m$. Then*

$$E[d(x_{j+1}, S)^2 \mid x_j] \;\leq\; \Big(1 - \frac{1}{m\bar{\gamma}^2}\Big) d(x_j, S)^2.$$

Proof First, note that by Inequality 2.5, the distance $\|x_j - \bar{x}\|$ is nonincreasing in j. Hence if x_0 is sufficiently close to $\bar{x}$, then x_j is as well for all $j \geq 0$. Then, again using Inequality 2.5 (applied to the set S_i), we have, for all points $x \in S \subset S_i$,

$$\|x_j - x\|^2 - \|x_j - P_{S_i}(x_j)\|^2 \geq \|P_{S_i}(x_j) - x\|^2.$$

Taking the minimum over $x \in S$, we deduce

$$d(x_j, S)^2 - \|x_j - P_{S_i}(x_j)\|^2 \geq d(P_{S_i}(x_j), S)^2.$$

Hence

$$\begin{aligned} E[d(x_{j+1}, S)^2 \mid x_j] &= \frac{1}{m}\sum_{i=1}^{m} d(P_{S_i}(x_j), S)^2 \\ &\leq \frac{1}{m}\sum_{i=1}^{m}[d(x_j, S)^2 - d(x_j, S_i)^2] \\ &= d(x_j, S)^2 - \frac{1}{m}\sum_{i=1}^{m} d(x_j, S_i)^2 \\ &= d(x_j, S)^2 - \frac{1}{m} d(0, \Phi(x_j))^2 \\ &\leq \Big(1 - \frac{1}{m\bar{\gamma}^2}\Big) d(x_j, S)^2, \end{aligned}$$

using the definition of metric regularity. □

Note that metric regularity at $\bar{x}$ for 0 is a slightly stronger assumption than actually necessary for this result. Specifically, the above result holds as long as Equation 5.2 holds for all x near $\bar{x}$ with $\bar{b} = 0$ fixed, as opposed to the above definition requiring it to hold for all b near $\bar{b}$ as well.

For a moment, let $m = 2$ and consider the sequence of iterates $\{x_j\}_{j\geq 0}$ generated by the randomized iterated projection algorithm. By idempotency of the projection operator, there's no benefit to projecting onto the same set in two consecutive iterations, so the subsequence consisting of different iterates corresponds exactly to that of the non-randomized iterated projection algorithm. In particular, if $x_j \in S_1$, then

$$d(P_{S_2}(x_j), S)^2 \leq d(x_j, S)^2 - d(x_j, S_2)^2 = d(x_j, S)^2 - [d(x_j, S_2)^2 + d(x_j, S_1)^2]$$

since $d(x_j, S_1) = 0$. This gives us the following corollary, which also follows through more standard deterministic arguments.

Corollary 5.5 *If Φ is metrically regular at $\bar{x}$ for 0 with regularity modulus γ and $\bar{\gamma}$ is larger than γ, then for x_0 sufficiently close to $\bar{x}$, the 2-set iterated projection algorithm is linearly convergent and*

$$d(x_{j+1}, S)^2 \leq \Big(1 - \frac{1}{\bar{\gamma}^2}\Big) d(x_j, S)^2.$$

Further, consider the following refined version of the m-set randomized algorithm. Suppose $x_0 \in S_1$ and $i_0 = 1$. Then for $j = 1, 2, \ldots$, let i_j be chosen uniformly at random from $\{1, \ldots, m\} \backslash \{i_{j-1}\}$ and $x_{j+1} = P_{S_{i_j}}(x_j)$. Then we obtain the following similar result.

Corollary 5.6 *If Φ is metrically regular at $\bar{x}$ for 0 with regularity modulus γ and $\bar{\gamma}$ is larger than γ, then for x_0 sufficiently close to $\bar{x}$, the refined m-set randomized iterated projection algorithm is linearly convergent in expectation and*

$$E[d(x_{j+1}, S)^2 \mid x_j,\ i_{j-1}] \leq \Big(1 - \frac{1}{(m-1)\bar{\gamma}^2}\Big) d(x_j, S)^2.$$

A simple but effective product space formulation by Pierra in [16] has the benefit of reducing the problem of finding a point in the intersection of finitely many sets to the problem of finding a point in the intersection of 2 sets. Using the notation above, we consider the closed set in the product space given by

$$T = S_1 \times S_2 \times \ldots \times S_m$$

and the subspace

$$L = \{Ax : x \in E\}$$

where the linear mapping $A : \mathbf{E} \to \mathbf{E}^m$ is defined by $Ax = (x, x, \ldots, x)$. Again, notice that $\bar{x} \in \cap_i S_i \Leftrightarrow (\bar{x}, \ldots, \bar{x}) \in T \cap L$. One interesting aspect of this formulation is that projections in the product space $\mathbf{E}^m$ relate back to projections in the original space $\mathbf{E}$ by

$$\begin{aligned} (z_1, \ldots, z_m) \in P_T(Ax) \quad &\Leftrightarrow \quad z_i \in P_{S_i}(x) \quad (i = 1, 2, \ldots, m) \\ (P_L(z_1, \ldots, z_m))_i \quad &= \quad \frac{1}{m}(z_1 + z_2 + \ldots + z_m) \quad (i = 1, \ldots, m) \end{aligned}$$

This formulation provides a nice analytical framework: we can use the above equivalence of projections to consider the *method of averaged projections* directly, defined as follows.

Algorithm 5.7 *Let $S_1, \ldots, S_m \subseteq E$ be nonempty closed convex sets. Let x_0 be an initial point. For $j = 1, 2, \ldots$, let*

$$x_{j+1} = \frac{1}{m} \sum_{i=1}^{m} P_{S_i}(x_j).$$

Simply put, at each iteration, the algorithm projects the current iterate onto each set individually and takes the average of those projections as the next iterate. In the product space formulation, this is equivalent to $x_{j+1} = P_L(P_T(x_j))$. Expanding on the work of Pierra in [16], additional convergence theory for this algorithm has been examined by Bauschke and Borwein in [2]. Under appropriate regularity conditions, the general idea is that convergence of the iterated projection algorithm for two sets implies convergence of the averaged projection algorithm for m sets. In a similar sense, we prove the following result in terms of randomized projections.

Theorem 5.8 *Suppose $S = \cap_{i=1}^m S_i$ is non-empty. If the randomized projection algorithm of Theorem 5.4 is linearly convergent in expectation with rate α, then so is Algorithm 5.7.*

Proof Let x_j be the current iterate, x_{j+1}^{AP} be the new iterate in the method of averaged projections and x_{j+1}^{RP} be the new iterate in the method of uniformly randomized projections. Then note that:

$$x_{j+1}^{AP} = \frac{1}{m}\sum_{i=1}^{m} P_{S_i}(x_j) = E[x_{j+1}^{RP}].$$

By convexity of the S_i's, it follows that

$$d(x_{j+1}^{AP}, S) = d(E[x_{j+1}^{RP}|x_j], S) \leq E[d(x_{j+1}^{RP}, S)|x_j] \leq \alpha d(x_j, S),$$

by Jensen's Inequality. □

Hence, the method of averaged projections converges no more slowly than the method of uniformly random projections. In particular, under the assumptions of Theorem 5.4, the method of averaged projections converges with rate no larger than $1 - \frac{1}{m\bar{\gamma}^2}$.

References

[1] H. Akaike. On a successive transformation of probability distribution and its application to the analysis of the optimum gradient method. *Annals of the Institute of Statistical Mathematics*, 11:1–16, 1959.

[2] H. H. Bauschke and J. M. Borwein. On the convergence of von Neumann's alternating projection algorithm for two sets. *Set-Valued Analysis*, 1:185–212, 1993.

[3] J. W. Demmel. On condition numbers and the distance to the nearest ill-posed problem. *Numerische Mathematik*, 51:251–289, 1987.

[4] J. W. Demmel. The probability that a numerical analysis problem is difficult. *Mathematics of Computation*, 50(182):449–480, 1988.

[5] F. Deutsch. *Best Approximation in Inner Product Spaces*. Springer-Verlag, New York, 2001.

[6] A.L. Dontchev, A.S. Lewis, and R.T. Rockafellar. The radius of metric regularity. *Transactions of the American Mathematical Society*, 355(2):493–517, 2003.

[7] C. Eckart and G. Young. The approximation of one matrix by another of low rank. *Psychometrika*, 1:211–218, 1936.

[8] G. Golub and C. van Loan. *Matrix Computations*. Johns Hopkins University Press, 1996.

[9] O. Güler, A.J. Hoffman, and U. Rothblum. Approximations to solutions to systems of linear inequalities. *SIAM J. Matrix Anal. Appl.*, 16(2):688–696, 1995.

[10] A.J. Hoffman. On approximate solutions of systems of linear inequalities. *J. Res. Nat. Bur. Stand.*, 49:263–265, 1952.

[11] R. Horn and C. Johnson. *Matrix Analysis*. Cambridge University Press, 1999.

[12] A.S. Lewis, D.R. Luke, and J. Malick. Local convergence for alternating and averaged nonconvex projections. *Foundations of Computational Mathematics*, in revision.

[13] W. Li. The sharp Lipschitz constants for feasible and optimal solutions of a perturbed linear program. *Linear Algebra and its Applications*, 187:15–40, 1993.

[14] K.F. Ng and X.Y. Zheng. Hoffman's least error bounds for linear inequalities. *Journal of Global Optimization*, 30:391–403, 2004.

[15] J.-S. Pang. Error bounds in mathematical programming. *Mathematical Programming*, 79:299–332, 1997.

[16] G. Pierra. Decomposition through formalization in a product space. *Mathematical Programming*, 28:96–115, 1984.

[17] A. Quarteroni, R. Sacco, and F. Saleri. *Numerical Mathematics*. Springer-Verlag, 2007.

[18] J. Renegar. Perturbation theory for linear programming. *Mathematical Programming*, 65:73–91, 1994.

[19] J. Renegar. Incorporating conditions measures into the complexity theory of linear programming. *SIAM J. Opt.*, 5(3):506–524, 1995.

[20] J. Renegar. Linear programming, complexity theory and elementary functional analysis. *Mathematical Programming*, 70:279–351, 1995.

[21] R.T. Rockafellar and R.J.-B. Wets. *Variational Analysis*. Springer, Berlin, 1998.

[22] T. Strohmer and R. Vershynin. A randomized Kaczmarz algorithm with exponential convergence. *Journal of Fourier Analysis and Applications*, to appear.